

CLAIMS

What is Claimed is:

1. A method for providing analytical business reports that are based on relational databases that include tables each with a plurality of data fields, comprising the steps of:

denormalizing said relational database;

inverting said relational database such that said data fields of said tables in said relational database are stored in separate files that contain a row number field and a single data field; and

compressing at least one of said files.

2. The method of claim 1 wherein said step of compressing is performed on files contain repeating data stored in successive rows.

3. The method of claim 2 wherein said files include a first file with a partition field and a second file with an analytical field.

4. The method of claim 3 further comprising the step of:
distributing processing of said business report by creating sub-rowsets of said second file and by assigning said sub-rowsets of said second file to a plurality of servers.

5. The method of claim 3 further comprising the step of:

distributing processing of said partition field by creating sub-rowsets of said partition field and by assigning said sub-rowsets of said partition field to a first set of servers, wherein said first set of servers identify unique partition values contained in said sub-rowsets.

6. The method of claim 5 further comprising the step of:

merging said unique partition values identified by said first set of servers;

transmitting said merged partition values to a second set of servers; and

processing said analytical field for said merged partition values by creating sub-rowsets of said analytical field and by assigning said sub-rowsets of said analytical field to said second set of servers

7. The method of claim 6 further comprising the steps of:

performing an analytical function using said second set of servers to create partial result sets; and

merging said partial result sets into a complete result set.

8. The method of claim 7 further comprising the step of:

annotating said result set.

9. The method of claim 8 further comprising the step of:

transmitting said complete result set to a client computer.

10. The method of claim 1 further comprising the step of:

directly traversing said at least one of said compressed files while said compressed file is stored in memory.

11. The method of claim 1 further comprising the step of:

receiving business user requests for said business reports via a browser of a computer that is connected to a distributed communication system.

12. The method of claim 11 further comprising the step of:

distributing said business user requests from multiple business users for said business reports over a plurality of web application servers using dynamic binding.

13. A system for providing analytical business reports, comprising:

a data storage device that stores denormalized and inverted data in separate files that contain a row number and a single data field, wherein a first file contains partition values and a second file contains analytical values;

a client computer that requests a business report that requires a calculation involving said first and second files;

an application controller that is responsive to said request from said client computer; and

a plurality of servers connected to said data storage device and said application controller,

wherein said application controller distributes portions of said calculation to said servers, merges partial result sets that are generated by said servers and transmits said merged result sets to said client computer.

14. The system of claim 13 wherein said first file is compressed.

15. The system of claim 13 wherein said first file is compressed using run length encoding.

16. The system of claim 13 wherein said application controller divides at least one of said first and second files into sub-rowsets and assigns said sub-rowsets to said servers for processing.

17. The system of claim 13 wherein said application controller distributes processing of said first file by creating sub-rowsets of said first file and by assigning said sub-rowsets of said first file to a first set of servers.

18. The system of claim 13 wherein said application controller distributes processing of said first file by creating sub-rowsets of said first file and by assigning said sub-rowsets of said first file to a first set of servers, wherein said first set of servers identify unique partition values in said sub-rowsets of said first file.

19. The system of claim 18 wherein said application controller merges said unique partition values identified by said first set of servers and transmits said merged partition values to a second set of servers.

20. The system of claim 19 wherein said second set of servers process sub-rowsets of said second file using said merged partition values to create partial result sets and wherein said second set of servers transmit said partial result sets to said application controller.

21. The system of claim 20 wherein said application controller merges said partial result sets into a complete result set.

22. The system of claim 21 wherein said application controller annotates said complete result set and transmits said complete result set to said client computer.

23. The system of claim 14 wherein said servers directly traverse said compressed file while said compressed file is stored in memory.

24. The system of claim 13 wherein said application controller receives said requests from said client computer via a browser run by said client computer.

25. The system of claim 13 further comprising the step of:
a web application server connected to a plurality of application controllers, wherein said web application server distributes said requests to said plurality of application controllers using dynamic binding.

26. The system of claim 13 wherein said application controller includes a distribution analysis module that determines a number of servers required to process said request within interactive time.

27. A method for providing analytical business reports that are based on relational databases that include tables each with a plurality of data fields, comprising the steps of:

denormalizing said relational database;

inverting said relational database such that said data fields of said tables in said relational database are stored in files that contain a row number and a data field;

compressing at least one of said files; and

distributing processing of said business report using a first set of servers.

28. The method of claim 27 further comprising the step of:

determining a number of servers in said first set that are required to process said request in interactive time.

29. The method of claim 27 wherein a first file includes a partition field and a second file includes an analytical field.

30. The method of claim 29 further comprising the step of:

dividing rows of said first file into sub-rowsets;

assigning said sub-rowsets to said first set of servers to determine unique partition values; and

merging said unique partition values generated by said first set of servers.

31. The method of claim 30 further comprising the step of:

distributing a control structure related to said unique partition values to a second set said servers;

dividing rows of said second file into sub-rowsets; and

assigning said sub-rowsets to said second set of servers.

32. The method of claim 31 further comprising the step of:

calculating partial result sets using said control structure, said analytical values and said unique partition values in said second set of servers.

33. The method of claim 32 further comprising the step of:

merging said partial result sets from said second set of servers into a complete result set;

annotating said complete result set; and

transmitting said complete result set to a client computer.

34. The method of claim 27 further comprising the step of:

directly traversing said compressed files while said compressed files are stored in memory.